SIAG/OPT Views-and-News

A Forum for the SIAM Activity Group on Optimization

Volume 10 Number 2 May 1999

Contents

Case Studies

An Application of Interior Methods to the Approximation of Forward Rate Curves

Anders Forsgren

Optimization and Systems Theory, Department of Mathematics, Royal Institute of Technology, SE-100 44 Stockholm, Sweden (andersf@math.kth.se). [1]

1. Introduction

The application considered in this article arises within finance and concerns approximation of a *forward rate* curve, given a set of *bond* prices. The simplest bond is a zero-coupon bond for which a given payment is obtained at a given time. This single payment may be normalized so that it is a unit payment. Consequently, for a given set of m zero-coupon bonds, it is known that bond i, $i = 1, \ldots, m$, gives unit value at time t_i, and in addition its market price at time zero, v_i, is known. Here, and throughout this article, the current time is set to zero and denoted by t_0. For a more general discussion, see, e.g., Björk [3, Chapter 15]. Without loss of generality, we assume that $0 = t_0 < t_1 < \cdots < t_m$. The purpose of the application is to extend these market prices to a smooth curve $v(t)$ which is consistent with the information that is known, i.e., $v(t_0) = 1$ and $v(t_i) = v_i$, $i = 1, \ldots, m$. Rather than working on $v(t)$, the *instantaneous forward rate* $f(t)$ is chosen. The instantaneous forward rate $f(t)$ is the riskless interest rate that can be obtained over an infinitesimal interval at time t for a decision made at time t_0. (See, e.g., Björk [3, Chapter 15] for a more precise definition.) This means that for $t \in [t_0, t_m]$, the relationship between $v(t)$ and the function f is given by

$$v(t) = e^{-\int_{t_0}^{t} f(s)ds}, \tag{1}$$

or equivalently

$$f(t) = -\frac{d}{dt} \ln v(t).$$

Since both f and v are unknown functions, the question is what functions are suitable as choices for f and v. One way of making such a choice is described below.

The background to the current article is that the author was approached by a Swedish company, Algorithmica Research AB, who wanted to utilize a certain model for approximation of forward rate curves, and implement this model in a tool to be used by financial analysts. The model is described in Section 2., the resulting optimization problem is

[1] Research supported by the Swedish Natural Science Research Council (NFR).

described in Section 3., and finally the resulting solution approach is discussed in Section 4.. This article gives a summary of Forsgren [4], where a more thorough discussion of the mathematical aspects is given.

2. Model

The chosen model, proposed by Adams and van Deventer [1], suggests a choice of f that gives "maximum smoothness" in the sense that the f that has minimal L_2-norm of its second-derivative over the planning horizon is chosen. This means that the objective is to minimize

$$\int_{t_0}^{t_m} (f''(s))^2 ds.$$

where it is required that f is continuously differentiable on the whole interval (t_0, t_m) and five times continuously differentiable on each subinterval (t_{i-1}, t_i), $i = 1, \dots, m$. Consistency with the market price v_i of bond i, for $i = 1, \dots, m$, gives the requirement

$$v_i = e^{-\int_{t_0}^{t_i} f(s)ds}, \quad i = 1, \dots, m,$$

utilizing (1). Equivalently, this may be written as

$$-\ln \frac{v_i}{v_{i-1}} = \int_{t_{i-1}}^{t_i} f(s)ds, \quad i = 1, \dots, m,$$

with $v_0 = 1$.

In addition, an extension suggested by Bjerksund and Stensland [2] is included, in which it is assumed that there are more general bonds such that, for each bond j, $j = 1, \dots, n$, today's bid price $\underline{p}^j$ and today's ask price $\bar{p}^j$ is known, in addition to the payments p_i^j at times t_i, $i \in I_j$, where I_j denotes the set of points of time where payments are made for bond j. The consistency requirement on v_i, $i = 1, \dots, m$, then becomes

$$\underline{p}^j \le \sum_{i \in I_j} p_i^j v_i \le \bar{p}^j, \quad j = 1, \dots, n.$$

For a more general discussion on modeling aspects, see, e.g., Adams and van Deventer [1], Bjerksund and Stensland [2], Frishling and Yamamura [6], Tanggaard [8], and the references given in these papers.

3. Optimization problem

The optimization problem resulting from the model described in Section 2. takes the form

$$\begin{array}{ll}
\underset{\substack{f \in C^1[t_0,t_m] \\ v \in I\!R^m}}{\text{minimize}} & \displaystyle\int_{t_0}^{t_m} (f''(s))^2 ds \\
\text{subject to} & \displaystyle\int_{t_{i-1}}^{t_i} f(s)ds = -\ln \frac{v_i}{v_{i-1}}, \quad i = 1, \dots, m, \\
& v_0 = 1, \\
& \underline{p}^j \le \displaystyle\sum_{i=1}^{m} p_i^j v_i \le \bar{p}^j, \quad j = 1, \dots, n, \\
& f \in C^5(t_{i-1}, t_i), \quad i = 1, \dots, m,
\end{array}$$

where m denotes the number of points of time where payments are made and n denotes the number of bonds, see Bjerksund and Stensland [2]. (Additional constraints that may be included, such as $f(t_0)$ being set to a known value, are omitted from the discussion here.)

Results on *natural splines*, see e.g., Schwarz [7, pp. 126–128], can be used to characterize the optimal f as a piecewise quadratic polynomial, which is three times continuously differentiable on $[t_0, t_m]$. For details, see Adams and van Deventer [1][2], Bjerksund and Stensland [2] or Forsgren [4]. This means that f may be modeled by piecewise quadratic polynomials, and that the unknowns are the coefficients in the polynomials.

If all bonds are zero-coupon bonds, i.e., $m = n$, and in addition the simpler model is used where ask price and bid price are set equal, then the values of v_i, $i = 1, \dots, m$, are uniquely determined. In this situation, the optimization problem is a convex equality-constrained quadratic-programming problem with $5m$ variables, the coefficients in the polynomials. An alternative to solving the optimality conditions of this quadratic-programming problem is to use the characterization of the natural splines, see the discussion in Forsgren [4].

If $m > n$, then the values of v_i are not uniquely determined, and the problem becomes nonlinearly constrained and in general nonconvex. Also in this situation, first-order optimality conditions can either be given by the "traditional" Karush-Kuhn-Tucker optimality conditions or by natural splines, see Forsgren [4].

[2] Corrections on the shape of the polynomials are given in van Deventer and Imai [9, Chapter 2].

4. Solution method

As discussed in Section 2., if all bonds are zero-coupon bonds, and in addition the ask price and bid price are set equal, then the resulting problem is a convex equality-constrained quadratic program whose solution can be obtained by solving the linear system of equations that the KKT-conditions form. Initial experiments confirmed that also if more general bonds were treated, but the ask price and bid price were set equal, the problem could be solved by a low number of Newton iterations on the first-order optimality conditions.

These findings led to an approach for solving the inequality-constrained problem, where an initial interior point was created by applying Newton's method to the first-order optimality conditions of the equality-constrained problem that results when the value of each bond is set to the mean value of its ask price and its bid price. Consequently, this initial problem takes the form

$$\begin{array}{ll} \underset{\substack{f\in C^1[t_0,t_m]\\ v\in I\!R^m}}{\text{minimize}} & \displaystyle\int_{t_0}^{t_m} (f''(s))^2 ds \\ \text{subject to} & \displaystyle\int_{t_{i-1}}^{t_i} f(s)ds = -\ln\frac{v_i}{v_{i-1}}, \quad i=1,\dots,m, \\ & v_0 = 1, \\ & \displaystyle\sum_{i=1}^{m} p_i^j v_i = \tfrac{1}{2}(\underline{p}^j + \bar{p}^j), \quad j=1,\dots,n, \\ & f\in C^5(t_{i-1},t_i), \quad i=1,\dots,m. \end{array}$$

Then a primal-dual interior method was applied to the solution of the original problem with this interior point as initial point. The primal-dual nonlinear equations were solved by Newton's method, see, e.g., Forsgren and Gill [5, Chapter 2.3]. The initial point which was created this way was found to be a good starting point, and rapid convergence was obtained. Based on these observations, although the problem may be nonconvex in general, no specific device was used to take care of nonconvexities, but the norm of the residual of the primal-dual equations was used as a merit function.

The implementation which was made by Algorithmica Research AB was written in C and C++. The heart of the computational effort was the solution of the Newton equations, which was carried out using a sparse matrix factorization routine. The linear equations that were solved this way became highly ill-conditioned as the solution was approached. This, however, did not seem to harm the convergence and acceptable response times were obtained. For the Swedish market, which the company had in mind, a typical problem size would be m approximately 60 and n approximately 20.

Acknowledgements

This application was initiated by Magnus Nyström at Algorithmica Research AB. The C/C++ implementation at Algorithmica Research AB was made by Karl-Johan Johnsson. Thanks are due to Tomas Björk and Camilla Landén, who have been very helpful in sharing their expert knowledge on financial mathematics.

REFERENCES

[1] K. J. Adams and D. R. van Deventer, *Fitting yield curves and forward rate curves with maximum smoothness*, Journal of Fixed Income, (1994), pp. 52–62.

[2] P. Bjerksund and G. Stensland, *Utledning av rentens terminstruktur ved "maksimum glatthet"-prinsippet*, Beta, (1996), pp. 2–6. ISSN 0801-3322 (In Norwegian).

[3] T. Björk, *Arbitrage Theory in Continuous Time*, Oxford University Press, 1998. ISBN 0-19-877518-0.

[4] A. Forsgren, *A note on maximum-smoothness approximation of forward interest rate*, Report TRITA-MAT-1998-OS3, Department of Mathematics, Royal Institute of Technology, Stockholm, Sweden, 1998.

[5] A. Forsgren and P. E. Gill, *Primal-dual interior methods for nonconvex nonlinear programming*, SIAM J. Optim., 8 (1998), pp. 1132–1152.

[6] V. Frishling and J. Yamamura, *Fitting a smooth forward rate curve to coupon instruments*, Journal of Fixed Income, (1996), pp. 97–103.

[7] H. R. Schwarz, *Numerical Analysis: A Comprehensive Introduction*, John Wiley and Sons, Inc., New York, 1989. ISBN 0-471-92064-9.

[8] C. Tanggaard, *Nonparametric smoothing of yield curves*, Review of Quantitative Finance and Accounting, 9 (1997), pp. 251–267.

[9] D. R. van Deventer and K. Imai, *Financial Risk Analytics: A Term Structure Model Approach for Banking, Insurance and Investment Management*, Irwin Professional, Chicago, IL, 1997. ISBN 0-7863-0964-4.

An Inverse Problem in Finance

Thomas F. Coleman
Department of Computer Science,
Cornell University, Ithaca, NY 14853. [1]

1. Introduction

Optimization principles underpin much of computational finance. Historically this has been particularly true in the area known as portfolio analysis - the selection and modification of a collection of financial instruments such as stocks, bonds, and options. Indeed, Harry Markowitz was a co-recipient of the 1990 Nobel Prize in economics in recognition of his work involving novel application of optimization concepts to portfolio analysis.

The 1997 Nobel Prize in economics was won by Robert Merton and Myron Scholes for their role in the discovery and development of the most famous equation in finance, the Black-Scholes (B-S) equation [2, 13]. This work has, at first blush, little to do with optimization. Instead, it rests on notions from stochastic calculus, differential equations, and finance 'arbitrage' principles. However, it should be noted that application of the Black-Scholes equation (and, more interestingly, a generalized version) requires knowledge of a mysterious parameter, the volatility. This is where numerical optimization can help.

In this article we discuss the (generalized) Black-Scholes equation and show how inverse optimization problems can be formulated to yield smooth volatility surfaces. The latter is useful both in the accurate pricing of exotic options as well as computing sensitivities.

2. The Black-Scholes (B-S) Equation

The B-S equation is the cornerstone of options pricing. Its' derivation is thoroughly covered in many introductory books on mathematical finance, e.g., [9, 10, 15]. Here we give a terse discussion: our objective is to quickly move to our main concern, the question of volatility and its computation via inverse optimization.

What is an option? An option is an example of a financial derivative. A financial derivative is a financial instrument that derives its value from another, i.e., from an *underlying*, such as a stock. The most basic option is the European vanilla option (sometimes just called a European option). The buyer of a European *put option* agrees to consider selling a share of the underlying at a predetermined strike price K and expiration time T. At time T this buyer has the option of *exercising*, i.e., selling the share at price K, or not. Typically the owner will sell if the current stock price is below K and will bypass the opportunity otherwise. A European *call option* is symmetric to the put option: the call option buyer agrees to consider buying a share of the underlying at time T at an agreed on strike price K. Having purchased the option for some amount of money, the buyer will typically exercise the option, i.e., purchase the share of the underying at price K and time T, if the current stock price at time T is greater than K. Otherwise, the option buyer will forego the opportunity.

Ignoring the cost of purchasing the option, the payoff on the put option is $P = K - S(T)$, if $K > S(T)$, and $P = 0$ otherwise (assuming that in the latter case the option is not exercised). More compactly, the payoff is $P = max(K - S(T), 0)$. Analogously, the buyer of a call option obtains a payoff of $P = max(S(T) - K, 0)$.

The vanilla European options are the simplest of the options. Many more complex options (often called exotic options) are also traded - we will consider only the vanilla European option in this article. One of the most important questions in finance is, what is a fair price of an option? The answer, under

[1] This expository article is targetted to a reader familiar with numerical optimization; little mathematical finance background is required. It can serve as an introduction to some of the basic ideas in options pricing as well as an illustration of an interesting application of numerical optimization.

common assumptions, is given by the B-S equation.

The challenge in pricing an option is that the underlying behaves in a stochastic way. The power of the B-S analysis is this: with a few assumptions about the behaviour of the underlying, and a knowledge of the standard deviation of the returns (or volatility), there is a rigorous and computable concept of a fair price. Moreover, this fair price is unique: under these same assumptions, an option that is priced differently from the B-S price is mispriced in that it will allow a trader to have a 'free lunch', i.e., make an immediate profit without risk. This is known as *arbitrage*, a condition which is typically assumed to be nonexistent[2] when developing a pricing model.

One of the basic assumptions behind the Black-Scholes model is that the behaviour of the underlying S follows geometric Brownian motion,

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t, \tag{2}$$

where the left-hand-side is an instantaneous relative change in the underlying (e.g., the stock value), μ is known as the 'drift' term (average rate of growth of the stock value) and μdt is the predictable component of the relative change in the stock price. The second term on the right-hand-side, σdW_t, corresponds to the random behaviour of the stock: σ is known as the *volatility* of the stock (volatility is the standard deviation, or square-root of the variance, of the returns of the underlying) and dW_t is random variable drawn from a normal distribution (the mean of dW_t is zero, the variance is dt).

Equation (2) has proven to be a very useful model of stock behaviour. Nevertheless, even with 'optimal' choices of σ and μ, it does not always capture reality. We discuss a generalization to overcome this gap in §2.

The question we are now faced with is how to value on option defined on an underlying S, e.g., stock, whose growth follows (2). The answer, under a number of assumptions including the no-arbitrage assumption mentioned above, satisfies the Black-Scholes partial differential equation:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 s^2 \frac{\partial^2 V}{\partial s^2} + rs\frac{\partial V}{\partial s} - rV = 0, \tag{3}$$

where V is the value of the option (i.e., that which we are trying to determine), r is the risk-free interest rate. Detailed derivation of (3) is given in most texts on mathematical finance, e.g., [9, 10, 15].

Interestingly, the drift term μ does not appear in equation (3). Consequently, knowledge of μ is not required to evaluate V. (This is known as risk neutral valuation.) However, additional information is required before (3) can be of practical value. In particular, boundary conditions are needed and a value for σ, the volatility, is required.

First, the boundary conditions: For both the European put and call options there are standard boundary conditions. For example, for the European call typical boundary conditions are:

$$\boxed{\begin{array}{l} \lim_{s\to+\infty} \frac{\partial V(s,t)}{\partial s} = e^{-(T-t)}, t \in [0,T] \\ V(0,t) = 0, t \in [0,T] \\ V(s,T) = max(s-K,0), \end{array}}$$

where T is the time to maturity, K is the strike price. The determination of volatility, σ, is the more interesting question. The methods for determining σ fall into two categories. First, σ can be estimated from first principles, i.e., based on the definition of standard deviation of the returns and using historical data. Typically this leads to 1-dimensional regression problem, e.g., [9, 15].

This first principles approach is easy to implement but has several unpleasant aspects. For example, there is the question of how far back in time to go and how the data should best be weighted. In addition, it is disconcerting that the regression solution for σ does not usually yield the actual price for any known traded option.

A second, more common approach, is calculation of *implied volatility* (implied vol) to yield a value for σ . Implied volatility is determined by solving a simple 1-dimensional inverse problem involving a similar traded option (with known price) on the same underlying. Implied volatility is thus that value of σ that, when substituted into equation (3) along with appropriate boundary conditions, yields the known price of the corresponding option. Thus for each

[2]In practise arbitrage opportunities do arise but typically disappear quickly due to the operation of opportunistic arbitrageurs.

traded option there is a corresponding implied vol, and one such value can be used to determine the fair price of a new option on the same underlying.

Use of implied vol is very common in the trading world. Indeed, implied vol is often 'quoted' instead of option prices. Every financial engineering software package includes an implied vol computation. For example, in the MATLAB Financial Toolbox [12], a Black-Scholes implied volatility computation for a European call option is invoked by

σ = **blsimpv(sc,K,r,T, call)**

where **sc** is the current price of the underlying, **K** is the strike price, **r** is the risk free interest rate, **T** is the time to maturity, and **call** is the value (or price) of the call option under consideration. Function **blsimpv** uses Newton's method.

Despite the popularity of the implied volatility concept, there are problems. In particular, it is unsatisfying that σ, in the derivation of (3) yields different values of σ for different options on the same underlying. An ensuing difficulty is, given several differing implied vol computations on the same underlying, how to choose a value for σ to price an exotic option on the same underlying. A more pernicious problem has to do with hedging. We refer the reader to basic books on finance, e.g., [9, 10, 15], for discussions of hedging strategies. Here it suffices to say that hedging involves computing the sensitivity of V with respect to different parameters. The choice of σ can greatly affect the derivative (and thus the hedging strategy) and so an arbitrary choice from a set of 'implied vols' can be misleading. Moreover, there is much evidence to indicate that σ varies with time and/or strike price (e.g., [5, 7, 8]): this suggests σ is better thought of as a function of (s,t), i.e., $\sigma = \sigma(s,t)$. A framework for this approach is discussed in §2.

Given the boundary conditions, there are several common ways to solve the B-S equation (3). First, with the aid of a look-up table to evaluate the normal cumulative distribution function, an explicit formula is available to determine the B-S value at an arbitrary stock value S. This is obviously very useful but suffers from the disadvantage that this approach does not generalize - most other types of options do not enjoy explicit solutions. Second, there are tree or lattice evaluators - these are commonly described in most introductory finance books. Such methods are easy to implement and do generalize to more complicated options. In fact, such methods can be viewed as explicit numerical approaches to the solution of (3). On the downside, explicit approaches can suffer from stability problems and can be costly (in an attempt to overcome the stability problems). Third, implicit PDE methods can be employed to solve (1.2). This approach also generalizes to more complex options, e.g., [14], but poses a problem for general path-dependent options. Finally, we remark that Monte-Carlo methods are heavily used to evaluate more complex options, especially where there is path-dependency. Monte-Carlo mthods are difficult to adapt to American-style options (i.e., options that allow the owner to exericise at any time before maturity).

3. Generalized Black-Scholes

A reasonable and realistic alternative is to think of the volatility as a surface, $\sigma = \sigma(s,t)$, rather than a constant. In particular, a more general model of the evolution of the stock price, replacing (3), is the 1-factor continuous diffusion approach:

$$\frac{dS_t}{S_t} = \mu(S_t,t)dt + \sigma(S_t,t)dW_t, \tag{4}$$

where both $\sigma = \sigma(s,t)$ and $\mu = \mu(s,t)$ are continuous differentiable functions of the underlying s and t. Note that S_t is a stochastic variable and W_t is standard Brownian motion. The value of a European option where the underlying is defined by (4) satisfies the generalized Black-Scholes equation [13]:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma(s,t)^2 s^2 \frac{\partial^2 V}{\partial s^2} + rs\frac{\partial V}{\partial s} - rV = 0. \tag{5}$$

Unlike the standard Black-Scholes equation, (5) does not enjoy an explicit solution; however, a discretized PDE approach can be used provided the surface $\sigma(s,t)$ is available for evaluation at all the grid points. Equation (5) obviously represents a potentially more powerful approach than the standard B-S equation which requires volatility to be a single number. But we are left with the question: how can the volatility surface be obtained?

Similar to the scalar case, an inverse (implied) point of view can be invoked. This approach uses current (or very recent) data, i.e., prices on recently traded options on the underlying of question. A straightforward implementation to this inverse problem, assuming model (5), yields a large-scale optimization problem.

To see this suppose we have m data triplets, $(\bar{v}_j, T_j, K_j)$, corresponding to recently traded options on the same underlying: the option price is $\bar{v}_j$, T_j is the time to maturity, and K_j is the strike price. Discretize s and t consistent with the numerical procedure to solve (5), to yield an M-vector s and N-vector t. Surface $\sigma(s,t)$ is thus represented as an M-by-N matrix Σ (of unknowns). We are now faced with a vastly underdetermined problem, in general, since the number of data points, m, satisfies $m << M * N$. (A typical value is $m = 20$ or $m = 30$ wheras the product $M * N$ could easily be $100,000$ or more). To take up the slack, and introduce smoothness into σ, Osher and Lagnado [11] propose minimization over $\sigma(s,t)$, Σ after discretization, of the function

$$\sum_{j=1}^{m}(v_j(\sigma(s,t))) - \bar{v}_j)^2 + \lambda\|\nabla\sigma(s,t)\|_2, \qquad (6)$$

where λ is a positive constant and $\|\cdot\|_2$ denotes the L^2 norm.

A difficulty with this approach, in addition to the delicate choice of the parameter λ, is the computational challenge. Problem (6) is a very large optimization problem ($M * N$ variables Σ). Moreover, the first term in (6) is very nonlinear and dense - density is due to the evaluation of the discretized aproximation of the PDE in (5) which depends on the entire surface $\sigma(s,t)$, i.e., each point of the matrix Σ is involved in the evaluation of (5) for each j, $j = 1, \cdots, m$.) Indeed, in light of the extreme computational expense, Osher and Lagnado compute only very approximate solutions to (6) using a steepest descent procedure. Unfortunately, approximate solutions yield rough surfaces σ. Rough volatility surfaces can cause severe pricing and hedging problems. An alternative approach [3] yields a smaller more tractable optimization problem. The solution σ is smooth. The essential idea is to build in smoothness from the start: assume $\sigma(s,t)$ is a bi-cubic spline, e.g., [1, 6], defined on $p \approx m$ knots. The knots are located in a regular way comensurate with the known data. In more detail, let the number of spline knots be $p \le m$. Choose a set of fixed spline knots $\{(\bar{s}_j, \bar{t}_j)\}_{j=1}^{p}$. Given the spline knots with corresponding local volatility values $\bar{\sigma}_i \stackrel{\text{def}}{=} \sigma(\bar{s}_i, \bar{t}_i)$, an interpolating cubic spline $c(s,t)$ with a fixed end condition (e.g., the natural spline end condition) is uniquely defined by setting $c(\bar{s}_i, \bar{t}_i) = \bar{\sigma}_i, i = 1, \cdots, p$. The freedom in this problem is represented by the volatility values $\{\bar{\sigma}_i\}$ at the given knots $\{(\bar{s}_i, \bar{t}_i)\}$. If $\bar{\sigma}$ is a p-vector, $\bar{\sigma} = (\bar{\sigma}_1, \cdots, \bar{\sigma}_p)^T$, then we denote the corresponding interpolating spline with the specified end condition as $c(s,t;\bar{\sigma})$. For $j = 1, \cdots, m$, let

$$v_j(c(s,t;\bar{\sigma})) \stackrel{\text{def}}{=} v(c(s,t;\bar{\sigma}), K_j, T_j).$$

To allow the possibility of incorporating additional a priori information, l and u are lower and upper bounds that can be imposed on the local volatilities at the knots. Thus, we define the *inverse spline local volatility approximation problem*: Given p spline knots, $(\bar{s}_1, \bar{t}_1) \cdots, (\bar{s}_p, \bar{t}_p)$, solve for the p-vector $\bar{\sigma}$

$$\min_{\bar{\sigma} \in I\!R^p} \quad f(\bar{\sigma}) \stackrel{\text{def}}{=} \tfrac{1}{2}\sum_{j=1}^{m}[v_j(c(s,t;\bar{\sigma})) - \bar{v}_j]^2 \qquad (7)$$

$$\text{subject to} \quad l \le \bar{\sigma} \le u. \qquad (8)$$

Note that (7) is a small optimization problem, typically with $m = p \approx 20$ variables, the solution has certain guaranteed smoothness properties (due to the use of the bi-cubic spline model), and, the given data will be usually be satisfied provided it is consistent.

4. Concluding Remarks

The volatility surface produced by the bi-cubic spline optimization approach discussed above is visually smooth in the area of interest [3]. Indeed, given the location of the knots it can be argued that the computed surface σ is the smoothest surface consistent with the (discretized) model (5) and the given data. However, the real test of any volatility surface computation, in addition to its computational attractiveness, is its useability vis-à-vis pricing and, especially, hedging with the generalized Black-Scholes model (5). Hedging involves computing the sensitivity of the option price with respect to different parameters. Experiments are currently in progress to determine the impact of using the bi-cubic surface representation for volatility in this context. (Efficient implementation of some sensitivity calculations involves applying either automatic differentiation or finite-differencing in a structured way [3, 4]).

We conclude with two remarks. First, while the bi-cubic spline optimization approach appears to produce a smooth, attractive, and useful volatility surface in the usual area of interest - in an $(s, t)-$region around known strike and expiration times for current option data - the volatility surface becomes less reliable outside of this region. This is usually not a problem but can be troublesome when pricing (or hedging with) long-dated options. This is an active area of investigation. Second, we expect that the bi-cubic optimization approach to this volatility surface problem can be applied to other inverse problems involving nonlinear, underdetermined systems.

5. Acknowledgements

Thanks to my colleagues Yohan Kim, Yuying Li, and Arun Verma for their help and suggestions.

REFERENCES

[1] G. Birkhoff and C. R. De Boor, *Piecewise polynomial interpolation and approximation*, in Approximation of Functions, H. L. Garabedian, ed., Elsevier (New York), 1965, pp. 164–190.

[2] F. Black and M. Scholes, *The pricing of options and corporate liabilities*, Journal of Political Economy, 81 (1973), pp. 637–659.

[3] T. F. Coleman, Y. Li, and A. Verma, *Reconstructing the unknown local volatility function*, The Journal of Computational Finance, to appear, 1999.

[4] T. F. Coleman and A. Verma, *Structure and efficient Jacobian calculation*, in Computational Differentiation: Techniques, Applications, and Tools, M. Berz, C. Bischof, G. Corliss, and A. Griewank, eds., SIAM, Philadelphia, Penn., 1996, pp. 149–159.

[5] S. R. Das and R. K. Sundaram, *Of smiles and smirks: A term-structure perspective*, Tech. Report NBER working paper No. 5976, Department of Finace, Harvard Business School, Harvard University, Boston, MA 02163, 1997.

[6] C. de Boor, *Spline Toolbox for use with MATLAB*, The Math Works Inc., 1997.

[7] E. Derman and I. Kani, *Riding on a smile*, Risk, 7 (1994), pp. 32–39.

[8] B. Dupire, *Pricing with a smile*, Risk, 7 (1994), pp. 18–20.

[9] J. Hull, Options, Futures, and Other Derivatives, 3rd Edition. Prentice Hall, 1997.

[10] R. Jarrow and S. Turnbull, Derivative Securities, South-Western Pub., 1996

[11] R. Lagnado and S. Osher, *Reconciling differences*, Risk, 10 (1997), pp. 79–83.

[12] Users Guide, *Financial Toolbox*, The Mathworks, Inc., Version 2, 1999.

[13] R. Merton, *The theory of rational option pricing*, Bell Journal of Economics and Management Science, 4 (1973), pp. 141–183.

[14] R. Zvan, K. Vetzal, and P. Forsyth, *Swing high Swing low*, Risk Magazine, (1998), pp. 71–74.

[15] P. Wilmott, and S. Howison, and J Dewynne, The Mathematics of Financial Derivatives. Cambridge University Press, 1997.

Valuation in Complete Markets: The Optimization Approach

Andrew Aziz
Algorithmics, Inc.

An optimization approach to valuation in complete markets that is based on a fundamental assumption of investor behaviour - that investors prefer more wealth to less. In market equilibrium, this single assumption is sufficient for valuation in an optimization framework.

Valuation methodologies in finance typically incorporate two fundamental assumptions of investor behaviour. The first is that investors prefer more wealth to less, and will take actions to maximize their wealth. The second is that investors have a distaste for risk and that distaste will lead them to modify their actions and make trade-offs between expected wealth and uncertainty of wealth.

The "preference of more to less" assumption is analytically quite tractable. It is intuitively very appealing and the notion of wealth maximization can be well defined quantitatively. In contrast, the "distaste of risk" assumption is much less tractable. It is somewhat less intuitively appealing and the notion of risk aversion can neither be well defined nor measured in a straightforward manner. As a consequence, valuation methodologies that are able to rely solely on the first assumption are typically much more successful than methodologies that must also rely upon the second assumption.

This tutorial provides an overview of a valuation framework within which the "preference of more to less" assumption can be quantitatively expressed as an optimization problem. In addition, we develop the necessary conditions under which the "preference of more to less" assumption is sufficient for valuation purposes in a frictionless market.

The optimization formulation

The behavioural assumption that investors prefer more to less guides the actions of arbitrageurs who seek to construct a position of holdings that maximize arbitrage profits. Arbitrage profits are typically classified on the basis of first- and second-order arbitrage. First-order arbitrage occurs when a portfolio can be constructed that yields positive cashflow today with no obligations in any future state. Second-order arbitrage occurs when a costless portfolio can be constructed that guarantees no future obligations but where there is a possibility of a strictly positive payoff in at least one future state.

The behaviour of arbitrageurs who act to maximize first-order arbitrage can be modeled as a multi-state, single period linear program:

$$\underset{\mathbf{x}}{\textbf{minimize:}} \quad c = p^T x \tag{1}$$

$$\textbf{subject to:} \quad A^T x \geq 0 \tag{2}$$

$$x \quad \text{unrestricted} \tag{3}$$

In this model, there are n independent securities and m possible future states. **A** represents the $n\,x\,m$ matrix of payoffs for security i in state j and a_{ij} is the guaranteed payoff of security i in state j. The market price that investors pay for security i is equal to p_i; **p** represents the $n\,x\,1$ vector of security prices. An investor purchases a portfolio with holdings in security i in the amount of x_i. It is assumed that there are no long or short constraints on the purchases of any of these securities. Thus, **x**, the $n \times 1$ vector of security holdings, is unconstrained. **0** represents an $m \times 1$ vector of zeros.

Case	Cash Flow		Linear Program Solution	
	Today	Future	Objective	Constraints
First-Order Arbitrage	cash extracted	no oblication	unbounded, $c = -\infty$	feasible
Second-Order Arbitrage	zero cost	no obligation, possibility of payoff	bounded, $c = 0$	$\sum_j a_{ij}x_i > 0$, for some j
Market Equilibrium	zero cost	no obligation, no payoff	bounded, $c = 0$	$\sum_j a_{ij}x_i = 0$, for all j

Table 1.1: Summary of arbitrage and equilibrium conditions

The formulation of this approach follows the multi-state, single period, discrete model developed by Ross [3], and extended by Dybvig and Ross [1] and by Prisman [2]. In their model, each possible future state, j, has a strictly positive probability of occurrence, $\pi_\mathbf{j}$. The vector π is an $m \times 1$ vector of state probabilities.

In this framework, an investor's objective to maximize arbitrage profits (Equation 1) is re-expressed as an objective to minimize the cost, c, of purchasing a portfolio. The state constraints (Equation 2) restrict the feasible solutions to those for which the net payoff over all trades is non-negative for each state. Investors are able to buy ($\mathbf{x} > \mathbf{0}$) or sell ($\mathbf{x} < \mathbf{0}$) in unrestricted quantities (Equation 3). Given that all investors observe the same payoff matrix, $\mathbf{A}$ (i.e., no differential taxation), and that there are no long and short constraints, the market is assumed to be frictionless.

When the value of the objective function is negative, cash has been extracted from the trades. When the value is zero the trades are costless. Since the "no trade" solution ($\mathbf{x} = \mathbf{0}$) is always feasible, the solution to the problem is non-positive.

Though the problem is feasible, it may be bounded or unbounded. There are three cases to consider. In the first case, the problem is unbounded and the solution to the objective function is negative infinity for a given set of security prices. This is the case of first-order arbitrage whereby unlimited cash may be obtained today with no future obligations. In the second and third cases the problem is bounded and the solution to the objective function is zero.

In the second case at least one of the constraints is solved as a strict inequality. This is the case of second-order arbitrage: given a costless initial investment, there is some possibility of a strictly positive future payoff. In the third case, all constraints are solved as strict equalities. This is known as the "no-arbitrage condition" since no first- or second-order arbitrage opportunities are available. The implications for the present and future cash flows, the objective function and constraints of the optimization solution are summarized in Table 1.

Market equilibrium

The no-arbitrage condition describes a market equilibrium whereby security prices adjust, in response to supply and demand forces, so as to eliminate any first- and second-order arbitrage opportunities. In other words, the actions of arbitrageurs, as they attempt to maximize arbitrage profits, serve to force prices of relatively underpriced securities higher and those of relatively overpriced securities lower until an equilibrium is reached. Price equilibrium exists when all first- and second-order arbitrage opportunities are eliminated.

The existence of efficient markets where investors can both observe and act on arbitrage opportunities, is critical to the assumption that observed security prices in market equilibrium are arbitrage free. The no-arbitrage condition is the key premise underlying most valuation methodologies used to price financial securities.

The dual of the problem

Given the absence of first-order arbitrage in market equilibrium (the primal problem is bounded), the following relationship may be derived from the dual problem:

$$Ad = p \tag{4}$$

and, as a consequence of the absence of second-order arbitrage (the state constraints are solved with equality):

$$d > 0 \tag{5}$$

where $\mathbf{d}$ represents the $m \times 1$ vector of dual prices and d_j the dual variable associated with state j. The vector of dual prices, $\mathbf{d}$, is also known as the "state price" vector. Equations 4 and 5 represent a statement of the "Fundamental Theorem of Asset Pricing", a formal proof of which is contained in Prisman (1986).

The dual variables have the standard interpretation. If an investor is required to earn a payoff of at least one unit in state j the objective function increases by an amount d_j. Accordingly, an arbitrary security, with a payoff of one unit in state j and zero otherwise, must have a price equal to d_j. Such an instrument is known as an Arrow-Debreu security. By definition, each Arrow-Debreu security is an independent security.

Complete markets

When the number of independent securities, n, is equal to the number of states, m, the market is said to be complete. Unique prices for all m (independent) Arrow-Debreu securities can only be determined in complete markets. A risk-free instrument has payoffs over the m states that are, by definition, equal. The price of a risk-free security paying one unit in each state must, therefore, be equal to the price of a portfolio containing one of each Arrow-Debreu security, j, or simply $\sum_j d_j$. The risk-free rate, r_f, associated with the term of the single period, must be:

$$\sum_{j=1}^{r_f=\frac{1}{m}-1} d_j \tag{6}$$

Furthermore, a security, i, with arbitrary payoffs can be expressed as a combination of m independent securities with holdings equal to the payoff in each state of the new security:

$$p_i = \sum_{j=1}^{m} a_{ij} d_j \tag{7}$$

Thus, any new security may be priced simply as a portfolio of Arrow-Debreu securities.

Note, significantly, that in a complete market, an arbitrary security can be priced without knowledge of either its expected payoff, $E\{a_{ij}\}$, or the probability of its payoffs, π_j, over each state. As a result, the pricing of securities in a complete market requires neither the explicit modeling of an investor's attitudes towards risk nor a description of the distribution of payoffs across states. Investor attitudes towards risk are entirely embedded in the prices of the Arrow-Debreu securities. As a consequence, only the assumption underlying the no-arbitrage condition, the "preference of more to less", is necessary for valuation purposes.

Risk-neutral valuation

The fact that state probabilities are unnecessary for valuation purposes motivates the approach of "risk-neutral valuation", so called because the valuations are independent of any assumptions made about investor attitudes towards risk. If investors are truly risk-neutral, the true expectation of any payoff distribution across states can be discounted at the risk-free rate. Since investors are not typically risk neutral, the expected payoffs cannot be discounted at the risk-free rate. Nonetheless, a simple algebraic manipulation of the dual relationship produces for each state, j, a new variable, π_j^*, that has the characteristics of a risk-neutral probability. The magnitude of each π_j^* is calculated such that the risk-neutral expected payoff across states, $\mathbf{A}\pi^*$, can be discounted at the risk-free rate to return the correct price, that is:

$$\left(\sum_{j=1}^{m} d_j\right) A\pi^* = p \tag{8}$$

where (rewriting Equation 6):

$$\frac{1}{1+r_f} = \sum_{j=1}^{m} d_j \tag{9}$$

and:

$$\sum_{j=1}^{\pi^*=\frac{d}{m}} d_j \tag{10}$$

The vector π^* represents the $m \times 1$ vector of risk-neutral probabilities. As the result stems from a simple algebraic manipulation, the risk-neutral approach is appropriate regardless of an investor's true attitude towards risk.

Two state, two security example

In a two-state, two-security market, the behaviour of arbitrageurs can be modeled as a linear programming problem:

$$\begin{aligned} \underset{x}{\text{minimize:}} \quad & c = p_1 x_1 + p_2 x_2 \\ \text{subject to:} \quad & a_{11} x_1 + a_{21} x_2 \geq 0 \\ & a_{12} x_1 + a_{22} x_2 \geq 0 \\ & x \text{ unrestricted} \end{aligned}$$

where p_1 represents the price of security 1, x_1 represents the holdings in security 1, and a_{12} represents the guaranteed payoff of security 1 in state 2.

Graphically, the primal problem can be represented by Figure 1:

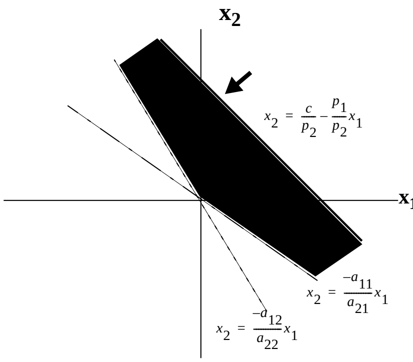

Figure 1: Feasible region and objective function

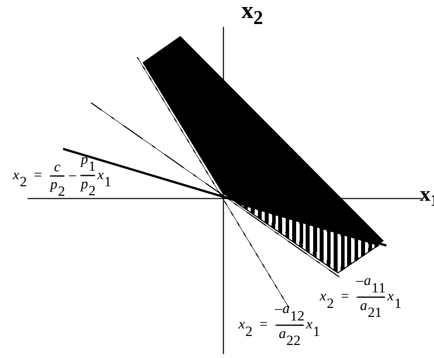

Figure 2: A case of first-order arbitrage

where the lines representing State Contraint 1 and State Constraint 2 can be represented, respectively, as:

$$x_2 = -\frac{a_{11}}{a_{21}} x_1$$

$$x_2 = -\frac{a_{12}}{a_{22}} x_1$$

and, where the line representing the objective function can be represented as:

$$x_2 = -\frac{p_1}{p_2} x_1$$

In the case of first-order arbitrage, the slope of the line representing the objective function is either greater than, or less than, the slopes of both of the lines representing the state constraints. Assume that the (absolute) slope of the objective function is less than the (absolute) slopes of the constraints (Figure 2).

Then:

$$\frac{a_{12}}{a_{22}} > \frac{a_{11}}{a_{21}} > \frac{p_1}{p_2}$$

The objective function and the constraints intersect at the origin, where $c = 0$. However, all holdings bounded by the objective function and State Constraint 1 that are long x_1 and short x_2 are feasible. As x_1 increases, c becomes more negative: the solution is unbounded and arbitrageurs reap unlimited first-order arbitrage until prices readjust.

If the slope of the line representing the objective function is equal to either of the slopes of the lines representing the state constraints, there are an infinite number of solutions and an opportunity for second-order arbitrage exists.

Assume that the objective function coincides with State Constraint 1 (Figure 3).
Then:

$$\frac{a_{12}}{a_{22}} > \frac{a_{11}}{a_{21}} > \frac{p_1}{p_2}$$

The objective function and the constraints intersect at the origin. At every point along the objective function $c = 0$. The distance between the objective function and State Constraint 2 is the slack in State Constraint 2. An arbitrageur is encouraged to increase his holdings (long x_1 and short x_2) to maximize the opportunity for payoff in state 2. Again, in market equilibrium, prices readjust until opportunities for second-order arbitrage are eliminated.

Recall that in no-arbitrage equilibrium, the relative prices of securities are such that the objective function has a solution equal to zero and the state constraints are solved as strict equalities. The no-arbitrage condition thus places bounds on relative security prices such that, at optimality, the objective function lies completely within the cone defined by the state constraints. In a two-state model, this implies that:

$$\frac{a_{12}}{a_{22}} > \frac{p_1}{p_2} > \frac{a_{11}}{a_{21}} \tag{11}$$

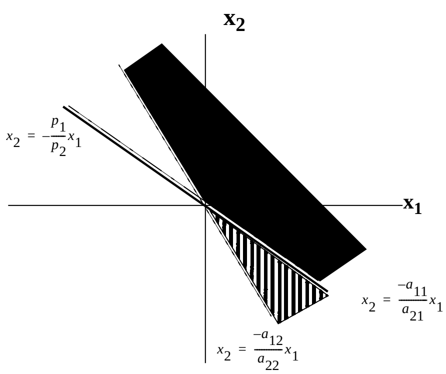

Figure 3: A case of second-order arbitrage

Again, the objective function and the constraints intersect at the origin (Figure 4).

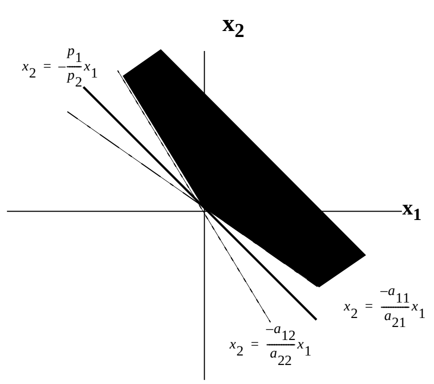

Figure 4: A case of market equilibrium

In market equilibrium, this is the unique, optimal solution. Any movement from the origin is either infeasible or sub-optimal.

Calculating dual variables

Given the no-arbitrage condition in market equilibrium, the following relationships may be derived from the dual problem (Equation 4):

$$a_{11}d_1 = a_{12}d_2 = p_1 \tag{12}$$

$$a_{21}d_1 = a_{22}d_2 = p_2 \tag{13}$$

where d_1 represents the price of the Arrow-Debreu security corresponding to State Constraint 1, a security that pays one unit in state 1 and zero otherwise. In this complete market ($n = m = 2$), unique values for the prices of Arrow-Debreu securities may be calculated as:

$$d_2 = \frac{p_2 - a_{21}d_1}{a_{22}} \tag{14}$$

$$d_1 = \frac{p_1 a_{22} - p_2 a_{12}}{a_{11}a_{22} - a_{12}a_{21}} \tag{15}$$

From Equation 7, the price of a new security, i, with payoffs, a_{i1} and a_{i2} over states 1 and 2 can be priced as a portfolio of Arrow-Debreu securities with holdings equal to the new security's payoffs in each state. Therefore, the equilibrium price of a new security, p_i, in market equilibrium is equal to:

$$p_i = a_{i1}d_1 + a_{i2}d_2 \tag{16}$$

Note that the risk-free security, with a payoff of one unit in each state, has a price equal to $d_1 + d_2$.

In risk-neutral valuation, the price of the arbitrary security with risk-neutral expected payoffs of $a_{i1}\pi_1^* + a_{i2}\pi_2^*$, is equal to the price of units of $a_{i1}\pi_1^* + a_{i2}\pi_2^*$ the risk-free security (Equation 8):

$$p_i = (a_{i1}\pi_1^* + a_{i2}\pi_2^*) * (d_1 + d_2) \tag{17}$$

where the risk-neutral probabilities, π_1^* and π_2^* for states 1 and 2, can be calculated as (Equation 10):

$$\pi_1^* = \frac{d_1}{d_1 + d_2} \tag{18}$$

$$\pi_2^* = \frac{d_2}{d_1 + d_2} \tag{19}$$

Summary

This tutorial provides an overview of the "Fundamental Theorem of Asset Pricing" as it applies to the valuation of securities in a complete and frictionless market. In the absence of arbitrage opportunities, a strictly positive and unique price exists for the Arrow-Debreu security associated with each state. This approach is developed within a multi-state single period model and is subsequently re-expressed in terms of a multi-period, single-state framework to model the bootstrapping problem.

However, in a market that is incomplete, the "preference of more to less" assumption is not sufficient to determine unique prices for the Arrow-Debreu securities, although bounds on these values may be found. As a consequence, valuation methodologies in incomplete markets require that further assumptions be made with respect to investor attitudes towards risk.

REFERENCES

[1] Dybvig, P. H. and S. A. Ross, *Tax clienteles and asset pricing*, Journal of Finance, 1986, 41(3): 751-771.

[2] Prisman, E. Z., *Valuation of risky assets in arbitrage free economies with frictions*, Journal of Finance, 1986, 41(3): 545-556.

[3] Ross, S., *Risk, return and arbitrage*, in Risk and Return in Finance, (I. Friend and J. Bicksler, eds.) Ballinger, Cambridge, MA, USA: 189-218, 1976.

Bulletin

Conferences

An International Conference on Complementarity Problems will be held in Madison, Wisconsin on June 9 - 12, 1999. The contemporary applications and algorithms that will be emphasized at the meeting will reflect the 35 years that have passed since complementarity was formally introduced and employed as a powerful mathematical model for a wide spectrum of problems in diverse fields.

The conference is intended to bring together engineers, economists, industrialists, and academicians from the U.S. and abroad who are involved in pure, applied, and/or computational research of complementarity and related problems.

The conference will consist of invited presentations, and is limited to 100 participants (including the speakers). A refereed volume of proceedings of the conference will be published. There are three major themes of the conference: engineering and machine learning applications, economic and financial applications, and computational methods. Each theme will be represented by several experts in the area.

Further details on the meeting, including registration deadlines, hotel and travel information can be found at: http://www.cs.wisc.edu/cpnet/iccp99

Michael Ferris
Olvi Mangasarian
Jong-Shi Pang
(co-organisers)

New Journal in Optimization and Engineering

Aims and Scope. Optimization and Engineering is a new multidisciplinary journal to be published by Kluwer Academic Publishers. Its primary goal is to promote the application of optimization methods in the general area of engineering sciences. This includes facilitating the development of advanced optimization methods for direct or indirect use in engineering sciences. The journal provides a forum in which engineering scientists obtain information about recent advances of optimization sciences, and researchers in mathematical optimization learn about the needs of engineering sciences and successful applications of optimization methods. Its aim is to close the gap between optimization theory and the practice of engineering.

All optimization methods of relevance to applications in engineering sciences will be considered: deterministic and stochastic, continuous, mixed integer and discrete, when they are relevant for applications in engineering sciences. The journal also strives to publish successful applications of optimization in various engineering areas.

The scope of the journal includes the following:

Optimization. All mathematical methods and algorithms of mathematical optimization. Numerical and implementation issues, optimization software, bench-marking, case studies. Specifically: linear and convex optimization, general nonlinear and nonlinear mixed-integer optimization, combinatorial optimization, equilibrium, multilevel and multiobjective optimization, stochastic optimization.

Engineering Sciences. Electrical engineering, VLSI design, robotics, mechanical and structural engineering, geophysical engineering, civil engineering, industrial engineering, chemical and process engineering, aerospace engineering, water management, environmental and bio-engineering, transportation and communication sciences.

SUBMISSIONS

Papers submitted for publication in OPTIMIZATION AND ENGINEERING should be sent to the editor in chief

Tamas Terlaky
Department of Statistics, Stochastic and Operations Research
Faculty of Information Technology and Systems
Delft University of Technology
P.O. Box 5031, 2600 GA Delft, The Netherlands
E-mail: T.Terlaky@twi.tudelft.nl
URL address: http://ssor.twi.tudelft.nl/ terlaky

Electronic submission of papers is strongly encouraged.

Comments from the Chair and Editor

This issue is focused entirely on optimization applied to the problems of finance. Our hope is that we will have one issue each year devoted to one particular application area. If you have any ideas that you'd like to suggest for application areas, please contact any SIAG/OPT officer.

We would like to remind you that this is your newsletter. Please consider submitting an expository article on your favorite optimization application (or any other interesting aspect of optimization.). Such articles can be very useful in spreading the word about optimization and its importance, throughout the SIAG/OPT community and beyond. They can also be wonderful classroom aids, helping students see the importance of optimization and applied mathematics more generally. Please contribute.

Thomas F. Coleman, SIAG/OPT Chair
Computer Science
Cornell University
Ithaca, NY 14853
coleman@tc.cornell.edu

Juan C. Meza, Editor
Sandia National Laboratories
P.O. Box 969, MS 9011
Livermore, CA 94551
meza@ca.sandia.gov